

The M4RIE library for dense linear algebra over small fields with even characteristic

Martin R. Albrecht
malb@lip6.fr

November 30, 2011

1 Introduction

Linear algebra over small finite fields has many direct applications, such as cryptography and coding theory. Other applications include efficient linear algebra over the rationals, e.g., by reducing such computations to a series of computations modulo small primes, and solving non-linear systems of equations using Gröbner bases [15]. For the latter recent work has emphasised that the right choice of linear algebra algorithms and implementations can make a significant impact on the performance of Gröbner basis algorithms [10].

Compared to other small finite fields, those with even characteristic have some special properties which make them a prominent choice for designing cryptographic and coding systems – cf., the AES [8] as a prime example. For instance, addition is simply XOR and is hence natively available on modern CPUs, the same cannot be said for other small finite fields. Yet, this family of finite fields has not received much attention in the literature on linear algebra. However, that the current state of the art in the literature leaves something to be desired can be observed from the following simple benchmark: Computing the (reduced) row echelon form of a random $4,000 \times 4,000$ matrix over \mathbb{F}_4 on a 2.66 Ghz Intel i7 CPU takes 5s using GAP 4.4.12 [11] (not reduced) or even 876.6s and 2805.8s using NTL 5.4.2 [17] (not reduced) and Sage [19] (reduced) respectively. Furthermore, LinBox’s FFPACK [9] performs the same task over the prime field \mathbb{F}_3 (reduced) in 4.55s. For comparison, the closed source system Magma [6] can compute the reduced row echelon form of a dense $4,000 \times 4,000$ matrix over \mathbb{F}_4 in 0.64s (reduced) and the same operation over \mathbb{F}_2 takes 0.054s using the M4RI library [3].

In this work, we present the M4RIE library which implements efficient algorithms for linear algebra with dense matrices over \mathbb{F}_{2^e} for $2 \leq e \leq 10$.¹ As the name of the library indicates, it makes heavy use of the M4RI library [1] both directly (i.e., by calling it) and indirectly (i.e., by using its concepts). The contributions of this work are as follows. We provide an open-source GPLv2+ C library for efficient linear algebra

¹Future versions will support $e \leq 16$.

over \mathbb{F}_{2^e} with $2 \leq e \leq 10$. In this library we implemented an idea due to Bradshaw and Boothby [5] which reduces matrix multiplication over \mathbb{F}_{p^n} to a series of matrix multiplications over \mathbb{F}_p . Furthermore, we propose a caching technique – *Newton-John* tables – to avoid finite field multiplications which is inspired by Kronrod’s method (“M4RM”) [4, 2] for matrix multiplication over \mathbb{F}_2 . Using these two techniques we provide asymptotically fast triangular solving with matrices (TRSM) and PLE-based [13] Gaussian elimination. As a result, we are able to significantly improve upon the state of the art in dense linear algebra over \mathbb{F}_{2^e} with $2 \leq e \leq 10$. For example, the above benchmark is completed in 0.4s by our library.

2 Notation

We represent elements in $\mathbb{F}_{2^e} \cong \mathbb{F}_2[x]/\langle f \rangle$, with $f \in \mathbb{F}_2[x]$, $\deg(f) = e$ and f irreducible, as polynomials $\sum_{i=0}^{e-1} a_i x^i$ or as coefficient vectors (a_{e-1}, \dots, a_0) where $a_i \in \mathbb{F}_2$. We sometimes identify the coefficient vector (a_{e-1}, \dots, a_0) with the integer $\sum_{i=0}^{e-1} a_i 2^i$, e.g., when indexing tables by finite field elements. By α we denote some root of the primitive polynomial f of \mathbb{F}_{2^e} . By A_i we denote the i -th row of the matrix A and by $A_{i,j}$ the entry in row i and column j of A . We start counting at zero. We represent permutation matrices as LAPACK-style permutation vectors. That is to say, that for example the permutation matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

is stored as $P = [0, 2, 2]$, where for each index i the entry P_i encodes which (row or column) swap should to be performed on the input matrix. This allows to apply permutations in-place.

3 Matrix representation

The M4RIE library features two matrix types, each of which is optimised for certain operations. Both representations use one or more M4RI matrices as data storage and hence re-use M4RI’s matrix window concept [2], allocation routines and data structures.

3.1 packed: mzed_t

Considering the polynomial representation of elements in \mathbb{F}_{2^e} , we may bit-pack several such elements in one machine word. Since we are using the M4RI library as actual data storage, this means that words hold 64 bits [2]. Hence, the `mzed_t` data type packs elements of \mathbb{F}_{2^e} in 64-bit words. However, instead of packing as many elements as possible into one word, every element is padded to the next length dividing 64. Thus, for example, elements in \mathbb{F}_{32} are represented as polynomials of degree 8 where the

top three coefficients are always zero. While this wastes some storage space and CPU time, it allows for more compact code by reducing what cases have to be considered. The second row of Figure 1 gives an example.

In this representation additions are very cheap since we can disregard any element boundaries and simply call M4RI's addition routines. Scalar multiplication, on the other hand, is much more expensive. Either we perform a table look-up for each *element* or we perform bit operations on words which perform multiplication and modular reduction in parallel on all elements of a word. In either case, multiplication is considerably more expensive than addition.

3.2 sliced: mzd_slice_t

Instead of representing matrices over \mathbb{F}_{2^e} as matrices over polynomials we may represent them as polynomials with matrix coefficients. That is, for each degree we store matrices over \mathbb{F}_2 which hold the coefficients for this degree. Hence, the data type `mzd_slice_t` for matrices over \mathbb{F}_{2^e} internally stores e -tuples of M4RI matrices, i.e., matrices over \mathbb{F}_2 . We call each M4RI matrix for some degree i a *slice* and refer to the operation converting from `mzed_t` to `mzd_slice_t` as *slicing*. The inverse operation is called *clinging*. The third row of Figure 1 gives an example of the `mzd_slice_t` representation.

Addition is performed by adding each slice independently and hence is quite efficient. Scalar multiplication, on the other hand, has to rely on similar techniques as in `mzed_t`. However, in addition, data locality in `mzd_slice_t` is worse than in `mzed_t`. Thus, here too, scalar multiplication is much more expensive than addition.

Thus, in this work, we present algorithms for matrix multiplication and elimination where we avoid many scalar multiplications.

$$\begin{aligned}
A &= \begin{pmatrix} \alpha^2 + 1 & \alpha \\ \alpha + \alpha + 1 & 1 \end{pmatrix} \\
&= \begin{bmatrix} \square 101 & \square 010 \\ \square 011 & \square 001 \end{bmatrix} \\
&= \left(\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \right)
\end{aligned}$$

Figure 1: 2×2 matrix over \mathbb{F}_8

4 Newton-John tables

To explain the main idea behind Newton-John tables, consider matrix multiplication as given in Algorithm 1.

Algorithm 1: Cubic matrix multiplication

Input: $A - m \times \ell$ matrix
Input: $B - \ell \times n$ matrix
Output: $C = A \cdot B$

```

1 begin
2   for  $0 \leq i < m$  do
3     for  $0 \leq j < \ell$  do
4        $C_j \leftarrow C_j + A_{j,i} \times B_i;$ 
5   return  $C;$ 

```

This algorithm uses $m \cdot \ell \cdot n$ finite field multiplications and the same number of additions. That is, in line 4 the row B_j is scaled by $A_{j,i}$ and then added to the row C_j . Observe that B_j is rescaled ℓ -times, while there are 2^e different values for $A_{j,i}$ and hence multiples of B_j . Indeed, if $2^e < \ell$ it is advantageous to precompute all possible 2^e multiples of B_j and to store these multiples in a table indexed by finite field elements. These precomputation tables are quite similar to Kronrod's method for matrix multiplication, also sometimes referred to as "greasing". Hence, we call these tables *Newton-John tables* to honour Olivia Newton-John's work in [14].

We also note that we create these tables in less than 2^e multiplications. That is, we first compute $\alpha^i \cdot B_j$ for all $0 \leq i < e$. Then, we compute each multiple of B_j as a linear combination of $(\alpha^0 \cdot B_j, \dots, \alpha^{e-1} \cdot B_j)$ which we just computed. Using Gray codes for the addition step we can thus construct all 2^e multiples using 2^e additions [12]. The subroutine creating these tables is given in Algorithm 2.

The complete algorithm is given in Algorithm 3 which costs $m \cdot (2^e + \ell) \cdot n$ additions and $m \cdot e \cdot n$ multiplications. Note that e is a constant here and asymptotically we thus

achieve $\mathcal{O}(n^3)$ additions and $\mathcal{O}(n^2)$ multiplications.

Algorithm 2: MAKETABLE

Input: B – an $1 \times n$ matrix
Output: T – a $2^e \times n$ matrix with each row a multiple of B

```

1 begin
2    $M \leftarrow e \times n$  matrix;
3    $T \leftarrow 2^e \times n$  matrix;
4   for  $0 \leq k < e$  do
5      $M_k \leftarrow \alpha^k \cdot B$ ;
6    $T \leftarrow$  all linear combinations of rows of  $M$ ;
7 return  $T$ ;
```

Algorithm 3: Newton-John multiplication

Input: A – $m \times \ell$ matrix
Input: B – $\ell \times n$ matrix
Output: $C = A \cdot B$

```

1 begin
2   for  $0 \leq i < m$  do
3      $T \leftarrow \text{MAKETABLE}(B_i)$ ;
4     for  $0 \leq j < \ell$  do
5        $x \leftarrow A_{j,i}$  as an integer;
6        $C_j \leftarrow C_j + T_x$ ;
7 return  $C$ ;
```

Many variants of this basic algorithm are possible. For instance, we may use more than one Newton-John table or process the data in blocks for better cache friendliness (cf., [2] for both techniques). Furthermore, if 2^e is too big to precompute T we may precompute only M (cf., Algorithm 2) and perform e additions in line 6. Our library uses eight Newton-John tables and processes matrices in blocks that fit into L2 cache. Since e is small we always compute the full table T .

Of course, this algorithm is not asymptotically fast. Hence, we only use it as a base case for the Strassen-Winograd algorithm [20] for matrix multiplication which has complexity $\mathcal{O}(n^{\log_2 7})$. In our implementation we cross over to the base case roughly when the submatrices fit into L2 cache; however, the exact value depends on the size of the finite field. Table 1 lists CPU times multiplying two $1,000 \times 1,000$ matrices in our implementation of Strassen-Winograd on top of Newton-John multiplication (abbreviated S-W/N-J in the following), in Magma and GAP (cf., Appendix A for a brief discussion of Magma’s and GAP’s implementations). Note that the hex string in the header of the last column indicates which revision of the public source code repository² was used to produce these times.

²cf., <https://bitbucket.org/m4rie>.

e	Magma 2.15-10	GAP 4.4.12	S-W/N-J 6b24b839a46f
2	0.013s	0.216s	0.012s
3	0.036s	0.592s	0.020s
4	0.074s	0.588s	0.022s
5	1.276s	1.568s	0.048s
6	1.286s	1.356s	0.059s
7	1.316s	1.276s	0.082s
8	1.842s	1.328s	0.160s
9	3.985s	64.700s	0.626s
10	4.160s	59.131s	1.080s

Table 1: Multiplication of $1,000 \times 1,000$ matrices on 2.66 Ghz Intel i7

4.1 Gaussian elimination

Newton-John tables can also be used in Gaussian elimination, as shown in Algorithm 4. This algorithm uses $r \cdot (n + 2^e) \cdot n$ additions and $r \cdot (e + 1) \cdot n$ multiplications, which gives an asymptotic complexity of $\mathcal{O}(n^3)$ additions and $\mathcal{O}(n^2)$ multiplications. Again, a variety of variants are possible such as multiple Newton-John tables (similar to [2]). Our implementation uses six Newton-John tables.

Algorithm 4: Newton-John Gauss elimination

Input: $A - m \times n$ matrix

Output: r – the rank of A

Result: A is in reduced row echelon form

```

1 begin
2    $r \leftarrow 0$ ;
3   for  $0 \leq j < n$  do
4     for  $s \leq i < m$  do
5       if  $A_{i,j} \neq 0$  then
6          $A_i \leftarrow A_{i,j}^{-1} \cdot A_i$ ;
7         swap the rows  $i$  and  $r$  in  $A$ ;
8          $T \leftarrow \text{MAKETABLE}(A_r)$ ;
9         for  $0 \leq k < m$  do
10          if  $k = s$  then continue;
11           $x \leftarrow A_{k,j}$  as an integer;
12           $A_k \leftarrow A_k + T_x$ ;
13         $r \leftarrow r + 1$ ;
14      break;
15 return  $r$ ;

```

In Table 2 we give CPU times for computing the reduced row echelon form of random $1,000 \times 1,000$ matrices over \mathbb{F}_{2^e} in Magma, GAP and our implementation. Note that GAP's `SemiEchelonMat` command does not compute the *reduced* row echelon form. Hence, to normalise the data we multiplied all GAP timings in Table 2 by two.

e	Magma 2.15-10	GAP 4.4.12	Newton-John 6b24b839a46f
2	0.028s	0.184s	0.012s
3	0.045s	0.496s	0.019s
4	0.054s	0.560s	0.022s
5	0.690s	1.224s	0.042s
6	0.670s	1.168s	0.048s
7	0.700s	1.104s	0.060s
8	0.866s	1.136s	0.081s
9	1.523s	35.634s	0.427s
10	1.540s	36.154s	0.831s

Table 2: Elimination of $1,000 \times 1,000$ matrices on 2.66 Ghz Intel i7

4.2 PLE decomposition

Algorithm 4 can be modified to compute the PLE decomposition instead of the row echelon form. Since this definition is lesser well-known we reproduce it below. For a more detailed treatment of PLE decomposition see [13].

Definition 1 (PLE). *Let A be a $m \times n$ matrix over a field K . A PLE decomposition of A is a triple of matrices P, L and E such that P is a $m \times m$ permutation matrix, L is a unit lower triangular matrix, and E is a $m \times n$ matrix in row-echelon form, and $A = PLE$.*

Lemma 1 ([13]). *A PLE decomposition of any $m \times n$ matrix A can be stored in-place in A .*

For the sake of simplicity, we compute a minor variant of PLE in our library. That is, L is not necessarily unit lower triangular, i.e., we do rescale the pivot row get leading entry 1. Then, the changes necessary for Algorithm 4 to compute this variant of PLE decomposition (in-place) are:

- Store i and j in two vectors P and Q in line 6;
- only start addition in column $c + 1$ in line 12 in order to keep L in place;
- perform column swaps below and on the main diagonal right before line 15 to compress L .

An alternative perspective on Newton-John-table based PLE decomposition is to consider it as a block iterative PLE decomposition (cf., [3]) with Newton-John table based multiplication updates to the right hand side.

Note, this algorithm is not asymptotically fast, hence its main application is as a base case for asymptotically fast PLE decomposition [13] which reduces PLE decomposition to matrix multiplication. For this, one last building block is needed:

4.3 TRIangular Solving with Matrices

Triangular system solving with matrices can also be achieved using Newton-John tables. As an example, we give an algorithm for solving $X = U^{-1} \cdot B$ with U upper triangular in Algorithm 5. We note that Algorithm 5 is essentially block iterative TRSM with Newton-John table based multiplication. Yet, we present it here for completeness.

Algorithm 5: Newton-John TRSM upper left

Input: U – $m \times m$ upper triangular matrix
Input: B – $m \times n$ matrix
Result: $X = U^{-1} \cdot B$ is stored in B

```

1 begin
2   for  $m > i \geq 0$  do
3      $B_i \leftarrow U_{i,i}^{-1} \cdot B_i$ ;
4      $T \leftarrow \text{MAKETABLE}(B_i)$ ;
5     for  $0 \leq j < i$  do
6        $x \leftarrow U_{j,i}$ ;
7        $B_j \leftarrow B_j + T_x$ ;

```

5 Karatsuba multiplication

Recall that `mzd_slice_t` represents matrices over \mathbb{F}_{2^e} as polynomials with matrices over \mathbb{F}_2 as coefficients. Using this representation, matrix multiplication then can be accomplished by performing polynomial multiplication and subsequent modular reduction. For example, assume we want to compute $C = A \cdot B$ where A and B are over \mathbb{F}_{2^2} . We rewrite A as $A_1x + A_0$ and B as $B_1x + B_0$, the product is then $\tilde{C} = A_1B_1x^2 + (A_1B_0 + A_0B_1)x + A_0B_0$ which reduces to $C = (A_1B_1 + A_1B_0 + A_0B_1)x + A_0B_0 + A_1B_1$ modulo the primitive polynomial $f = x^2 + x + 1$ of \mathbb{F}_{2^2} . Hence, matrix multiplication over \mathbb{F}_{2^e} can be reduced to matrix multiplication and addition over \mathbb{F}_2 . Using naive polynomial arithmetic we get that matrix multiplication over \mathbb{F}_{2^e} costs e^2 matrix multiplications over \mathbb{F}_2 . However, using Karatsuba polynomial multiplication we can reduce this to $e^{\log_2 3} \approx e^{1.584}$. To get back to the above example, we can rewrite it as $C = ((A_1 + A_0)(B_1 + B_0) + A_0B_0)x + A_0B_0 + A_1B_1$ and hence multiplication costs 3 instead of 4 multiplications over \mathbb{F}_2 . This was first explicitly proposed for matrices over \mathbb{F}_{p^n} by Bradshaw and Boothby in [5]. However, this technique has been used for linear algebra over \mathbb{F}_{2^k} with $2 \leq k \leq 4$ in Magma for some time [18].

Concrete costs are given in Table 3 where the first column lists the CPU time for multiplying two $4,000 \times 4,000$ matrices using Strassen-Winograd on top of Newton-John multiplication, the column “ M ” indicates how many $4,000 \times 4,000$ matrix multiplications over \mathbb{F}_2 can be achieved in the same time using the M4RI library, the column “naive” lists how many multiplications would be needed by naive polynomial multiplication, the column “[16]” lists the best known complexity for Karatsuba-like formulas, the last column shows the number of multiplications which our Karatsuba-like imple-

mentation actually achieves. The absolute time of our bitsliced implementation is given in the column “Bitslice”. Table 3 also compares our implementation with the previous two best implementations GAP and Magma .

However, we note that Karatsuba based multiplication needs more memory than Strassen on top of Newton-John multiplication. Our implementation uses three temporary matrices over \mathbb{F}_2 . We finish this section by pointing out in principle more efficient polynomial multiplication algorithms for $\mathbb{F}_2[x]$ can be applied (cf., [7]). However, due to the small degrees considered in this work it does not seem advantageous.

e	Magma 2.15-10	GAP 4.4.12	S-W/N-J	Bitslice	M	naive	[16]	Bitslice/ M4RI
2	1.220s	12.501s	0.630s	0.224s	8.8	4	3	3.1
3	2.020s	35.986s	1.480s	0.448s	20.8	9	6	6.3
4	5.630s	39.330s	1.644s	0.693s	23.1	16	9	9.7
5	94.740s	86.517s	3.766s	1.005s	53.0	25	13	14.2
6	89.800s	85.525s	4.339s	1.336s	61.1	36	17	18.8
7	82.770s	83.597s	6.627s	1.639s	93.3	49	22	23.1
8	104.680s	83.802s	10.170s	2.140s	143.2	64	27	30.1

Table 3: Multiplication of $4,000 \times 4,000$ matrices over \mathbb{F}_{2^e} on 2.66 Ghz Intel i7.

6 Echelon Forms

Putting these building blocks together

- (1) Karatsuba multiplication,
- (2) Newton-John-based PLE decomposition,
- (3) asymptotically-fast PLE decomposition,
- (4) Newton-John-based Triangular Solving with Matrices (TRSM) and
- (5) asymptotically-fast TRSM,

we can construct asymptotically fast Gaussian elimination, e.g., computation of (reduced) row echelon forms (cf., [13]). Our implementation uses `mzd_slice_t` as representation for large matrices and switches over to `mzed_t` when the submatrix currently considered fits into L2 cache. Table 4 lists CPU times for computing the (reduced) row echelon form using Magma (reduced), GAP (not reduced) and our implementation (reduced). Note that our implementation as of now only implements asymptotically fast PLE decomposition up to $e = 8$, for $e \in \{9, 10\}$ Newton-John table based Gaussian elimination is used.

e	Magma 2.15-10	GAP 4.4.12	M4RIE 6b24b839a46f
2	6.040s	162.658s	3.310s
3	14.470s	442.522s	5.332s
4	60.370s	502.672s	6.330s
5	659.030s	N/A	10.511s
6	685.460s	N/A	13.078s
7	671.880s	N/A	17.285s
8	840.220s	N/A	20.247s
9	1630.380s	N/A	260.774s
10	1631.350s	N/A	291.298s

Table 4: Elimination of $10,000 \times 10,000$ matrices on 2.66 Ghz Intel i7

References

- [1] Martin Albrecht and Gregory V. Bard. *The M4RI Library – Version 20111004*. The M4RI Team, 2011. <http://m4ri.sagemath.org>.
- [2] Martin Albrecht, Gregory V. Bard, and William Hart. Algorithm 898: Efficient multiplication of dense matrices over GF(2). *ACM Transactions on Mathematical Software*, 37(1):14 pages, January 2010. pre-print available at <http://arxiv.org/abs/0811.1714>.
- [3] Martin R. Albrecht, Gregory V. Bard, and Clément Pernet. Efficient dense Gaussian elimination over the field with two elements. *arXiv:1111.6549 [cs.MS]*, 2011.
- [4] V. Arlazarov, E. Dinic, M. Kronrod, and I. Faradzev. On economical construction of the transitive closure of a directed graph. *Dokl. Akad. Nauk.*, 194(11), 1970. (in Russian), English Translation in Soviet Math Dokl.
- [5] Tomas J. Boothby and Robert Bradshaw. Bitslicing and the Method of Four Russians over larger finite fields. *CoRR*, abs/0901.1413, 2009.
- [6] Wieb Bosma, John Cannon, and Catherine Playoust. The MAGMA Algebra System I: The User Language. In *Journal of Symbolic Computation* 24, pages 235–265. Academic Press, 1997.
- [7] Richard P. Brent, Pierrick Gaudry, Emmanuel Thomé, and Paul Zimmermann. Faster multiplication in GF(2)[x]. In *Proceedings of the 8th International Conference on Algorithmic Number Theory, ANTS-VIII’08*, pages 153–166, Berlin, Heidelberg, 2008. Springer-Verlag.
- [8] Joan Daemen and Vincent Rijmen. *The design of Rijndael: AES - the Advanced Encryption Standard*. Springer Verlag, Berlin, Heidelberg, New York, 2002.
- [9] Jean-Guillaume Dumas, Pascal Giorgi, and Clément Pernet. FFPACK: finite field linear algebra package. In *Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation, ISSAC ’04*, pages 119–126, New York, NY, USA, 2004. ACM.

- [10] Jean-Charles Faugère and Sylvain Lachartre. Parallel Gaussian Elimination for Gröbner bases computations in finite fields. In *Proceedings of the 4th International Workshop on Parallel and Symbolic Computation*, pages 89–97, 2010.
- [11] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4.12*, 2007.
- [12] Frank Gray. Pulse code communication, March 1953. US Patent No. 2,632,058.
- [13] Claude-Pierre Jeannerod, Clément Pernet, and Arne Storjohann. Fast gaussian elimination and the PLE decomposition. *in preparation*, page 30 pages, 2011.
- [14] Randal Kleiser, John Travolta, Olivia Newton-John, Jim Jacobs, Warren Casey, Bronte Woodard, and Allan Carr. *Grease*. Paramount Pictures, 1978.
- [15] Daniel Lazard. Gröbner-bases, Gaussian elimination and resolution of systems of algebraic equations. In *Proceedings of the European Computer Algebra Conference on Computer Algebra*, volume 162 of *Lecture Notes in Computer Science*, Berlin, Heidelberg, New York, 1983. Springer Verlag.
- [16] Peter L. Montgomery. Five, six, and seven-term Karatsuba-like formulae. *IEEE Trans. on Computers*, 53(3):362–369, 2005.
- [17] Victor Shoup. NTL. <http://www.shoup.net/ntl/>, 2009. version 5.4.2.
- [18] Allan Steel. Private communication. 29. November, 2011.
- [19] W. A. Stein et al. *Sage Mathematics Software (Version 4.7.1)*. The Sage Development Team, 2011. <http://www.sagemath.org>.
- [20] Volker Strassen. Gaussian elimination is not optimal. *Nummerische Mathematik*, 13:354–256, 1969.

A Other implementations

GAP [11] packs finite field elements of size $2 < s \leq 2^8$ into words using 8 bits per entry. Arithmetic is implemented using table look ups. Multiplication is performed using cubic matrix multiplication. Row echelon forms are computed using cubic Gaussian elimination.

LinBox/FFPACK [9] uses floating point numbers to represent finite field elements. For extension fields, elements are represented as “sparse” integers, such that there are enough zeroes between two coefficients to avoid the carry travelling too far. However, this feature is not readily exposed to the end-user and requires some tweaking to work. FFPACK implements Strassen-Winograd multiplication and asymptotically fast PLUQ decomposition for Gaussian elimination.

Magma [6] implements asymptotically fast matrix multiplication and reduces Gaussian elimination to LQUP decomposition. For \mathbb{F}_{2^k} with $2 \leq k \leq 4$ a bit-sliced representation similar to our `mzd_slice_t` is used in combination with Karatsuba-like formulas for polynomial multiplication. For $5 \leq k \leq 20$ elements in \mathbb{F}_{2^k} are represented using Zech logarithms. For larger k a packed polynomial representation is used similar to our `mzed_t` [18].